

We're pleased to welcome you to InfiniDB and infinidb.org ! Today, Calpont is making available the open source version of our analytic database, which we've named InfiniDB. InfiniDB is designed to handle analytic applications, data warehouses and marts, read-intensive applications, and read shards of large Web applications. To do that effectively, we've tried our best to provide a feature set that's capable of handling medium to very large data volumes and giving you fast response times at the same time. Right now, the core feature set of the InfiniDB Community Edition consists of the following:

- Column-oriented architecture
- Multi-threaded design to use all CPU's/cores on a box for a single query or load
- Transactional with crash recovery
- No need for indexing
- Vertical and automatic horizontal range partitioning of data
- High-speed data loader
- DML (insert, update, delete) support
- MVCC or 'snapshot read' support
- ALTER TABLE support (add,drop)
- Hash join support
- Logical data compression - SQL query diagnostics - No limits on RAM, CPU, size of database, or concurrent users
- MySQL front end

For more information on the above features and more on how InfiniDB works, check out our [Concepts Guide](#), which is part of our [free documentation set](#).

Please browse around the new [infinidb.org site](http://infinidb.org) . You'll be able to [download and use InfiniDB](#) (both binaries and source code), participate in the [forum](#), catch up on the latest news, surf through our wiki's, blogs, and other [resources](#)

, download and read our full set of [documentation](#), and contribute to the InfiniDB project.

Our goal with InfiniDB is to provide an analytic database that everyone can use; one that is open, fast, and simple to use. We hope you enjoy InfiniDB and we welcome your feedback!